Appendix A: Supplementary results

Figure 6 shows the complete state-by-state results for STV-based elections where the high-level patterns from Figure 1 can be seen in more granular detail. For example, the bump structure at uniform district sizes is clear for Michigan (MI), Florida (FL), Wisconsin (WI), California (CA), North Carolina (NC), and Pennsylvania (PA), among others. In CA, MA, NY, OK, and TN single-member proportional plans are not even possible due to the diffusion of minority party voters. In all states, the maximum gerrymandering capability peaks at about 1.5 seats per district and starts to decay rapidly with larger districts.

In Figure 7 we see that within our ensemble of maps, under full control of the redistricting process, Democrats could actually more effectively gerrymander the House of Representatives than Republicans, a finding counter to the conventional wisdom that geography favors Republican gerrymandering (Chen and Rodden 2013, Borodin et al. 2022). One difference may be that our analysis is still using only “natural” districts – those which aren’t wildly contorted – and therefore does not include the most surgical gerrymanders. This is a standard challenge in redistricting algorithms, which are effective for large-scale analysis but can often be outperformed (with respect to any metric) in any specific setting by experts. Regardless, as shown in Figure 6, it is true that Democrats can gerrymander their large states (CA, NY, and IL) more effectively than Republicans can gerrymander theirs (TX and FL), and that when ignoring the VRA, Democrats can crack large cities into many Democratic-leaning wedge-shaped districts. Furthermore, the gap in advantage is largest at two- and three-member districts because these thresholds enable Democrats to more efficiently place their voters than Republicans. For STV in particular, Democrats can break the $\frac{2}{3}$ or $\frac{3}{4}$ threshold needed for a sweep in many urban districts, while also still clearing the $\frac{1}{3}$ or $\frac{1}{4}$ threshold needed to gain one seat in rural districts, resulting in fewer wasted votes across most types of districts. While critics might point to this as a deal-breaker, it is important to recognize that this is only true in the limit of Democratic control, and such a scenario is deeply unrealistic.

In Figure 8, we show how the median, max Republican, and max Democratic absolute proportionality gap changes as a function of the voting rule and ratio of districts to seats. The median gap drops by about a factor of three between single-member and two-member districts and continues to slowly decay with larger districts. The median gap is a relevant metric because it proxies how easy it is to create a proportional map. Similarly, if the median map is fair, then there exist many fair maps, and it becomes easier to optimize for other desirable criteria like proportional racial representation, maintaining political subdivisions, and compactness. Also in Figure 8, we see that STV and Thiele squared track each other very closely, with the exception of Democratic gerrymanders of two- and three-member districts. This follows similar logic as the overall Democratic advantage analysis above, except that because Thiele squared requires more votes for a sweep, Democrats can no longer rely on full control of districts with just above $\frac{2}{3}$ of votes, and so end up wasting many votes in more heavily democratic areas by just missing the sweep threshold.

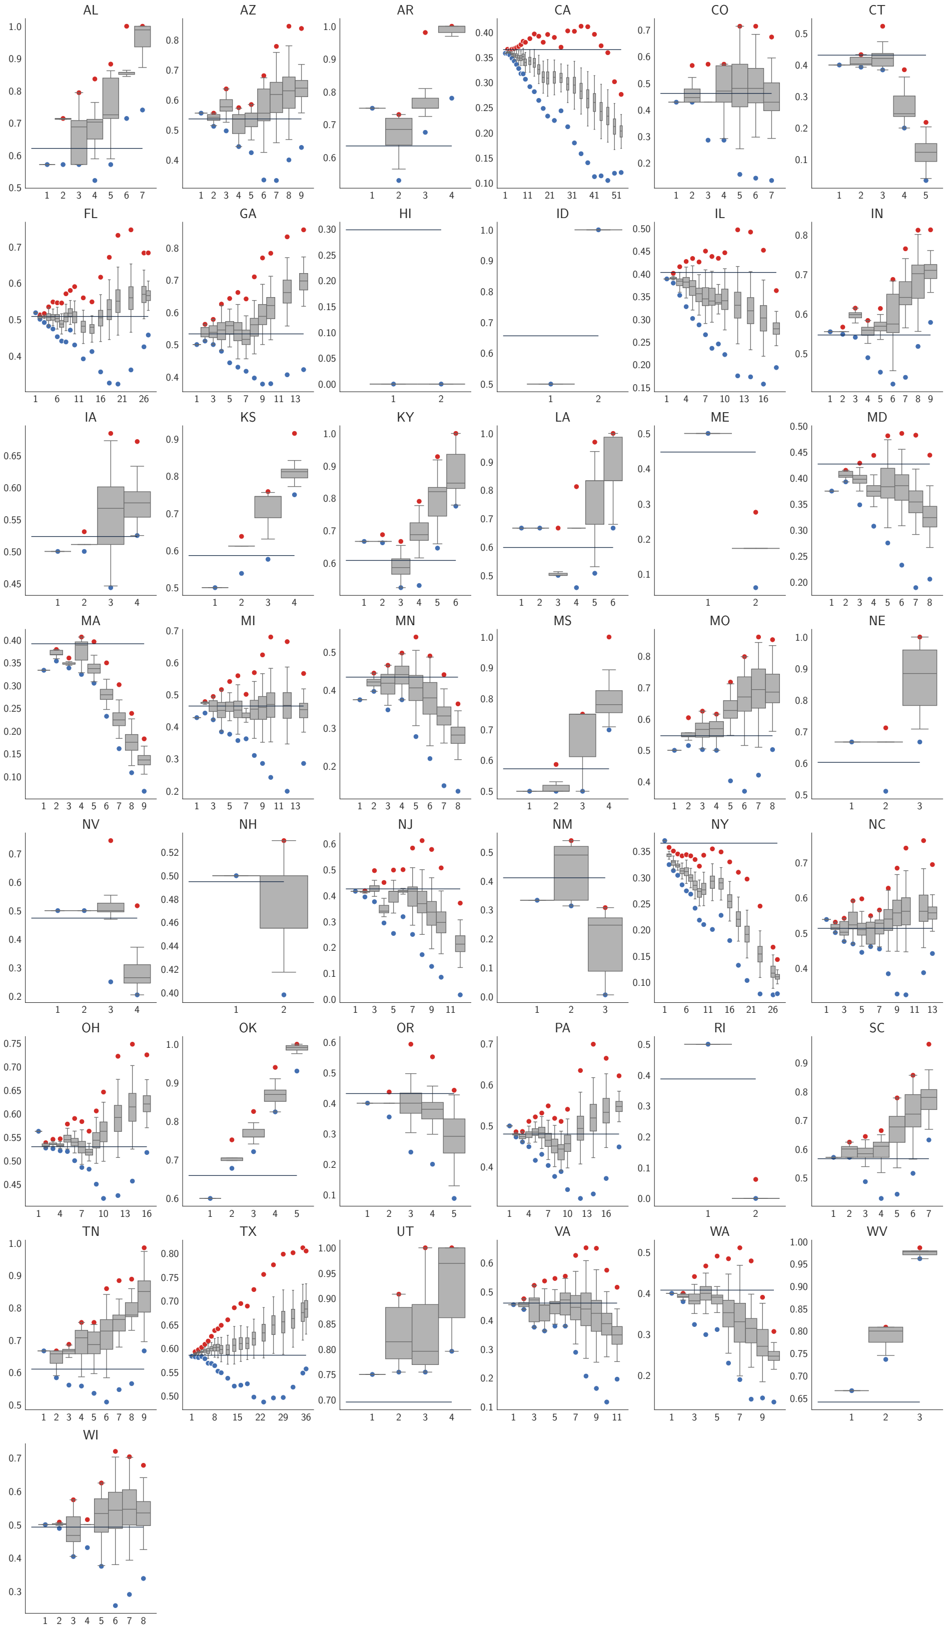

Figure 6 Figure 2c, repeated for each state with at least 2 Representatives. The y-axis is Republican seat share.

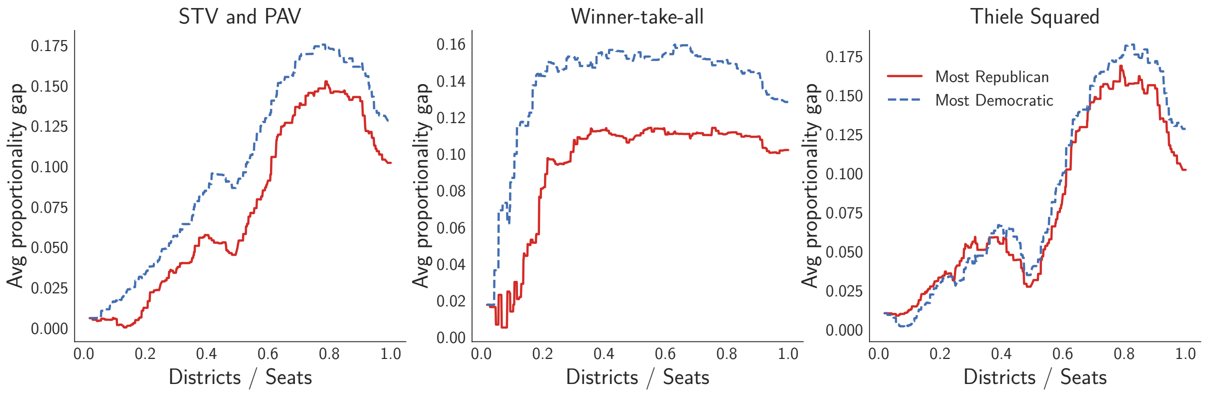

Figure 7 Proportionality gap in favor of each party in the most advantageous map

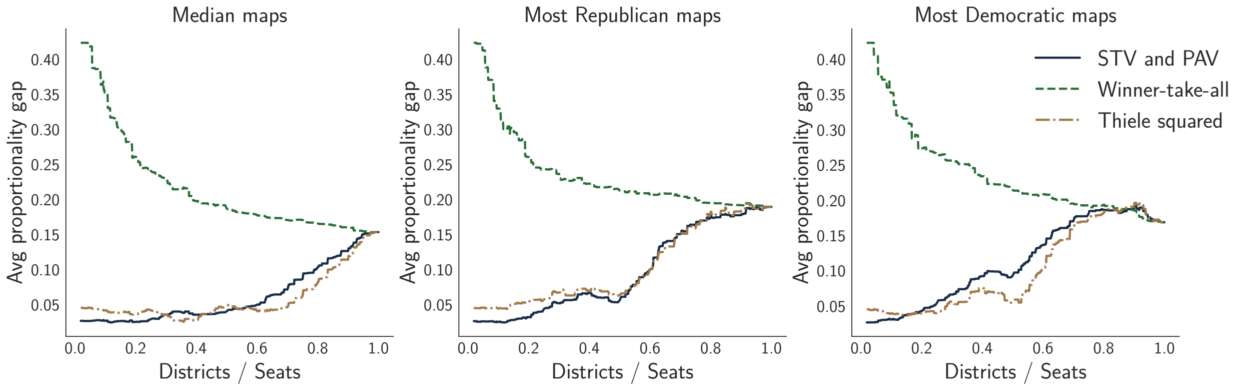

Figure 8 Figure 2a, repeated for Median, Most Republican, and Most Democratic maps.

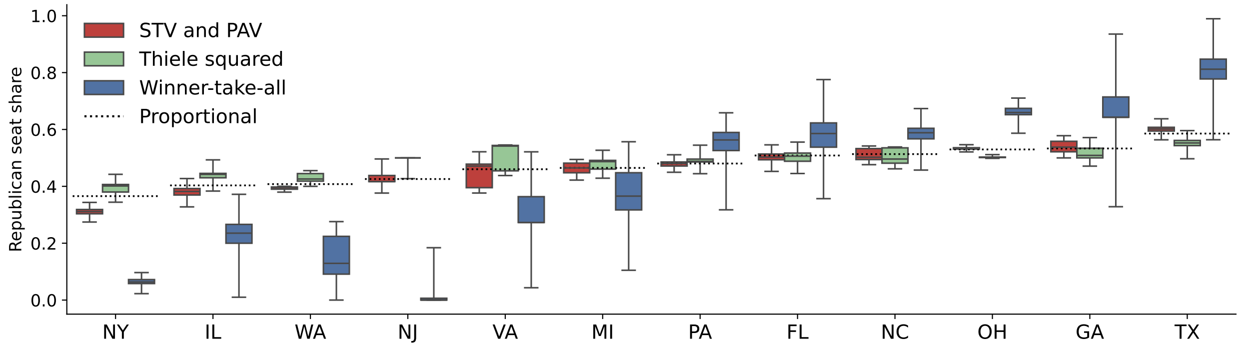

Figure 9 Seat share distribution of Fair Representation Act-compliant ensembles by voting rule.

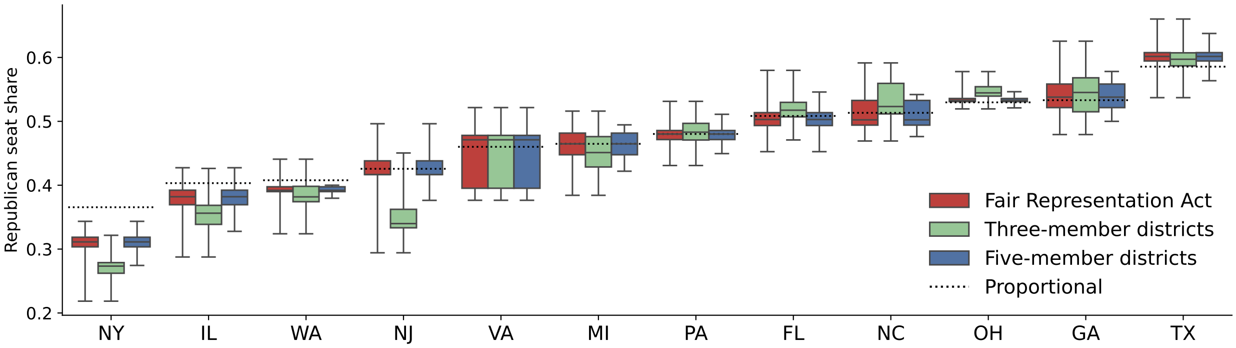

Figure 10 Comparison of Republican seat share distribution of Fair Representation Act-compliant ensembles with baseline for a selection of states using the STV/PAV voting rule (see Appendix Figure 9 for more rules). The Fair Representation Act has $N_k \in \{3,4,5\}$. The baseline is to use only three-member and only five-member districts (with four-member districts for overflow as necessary).

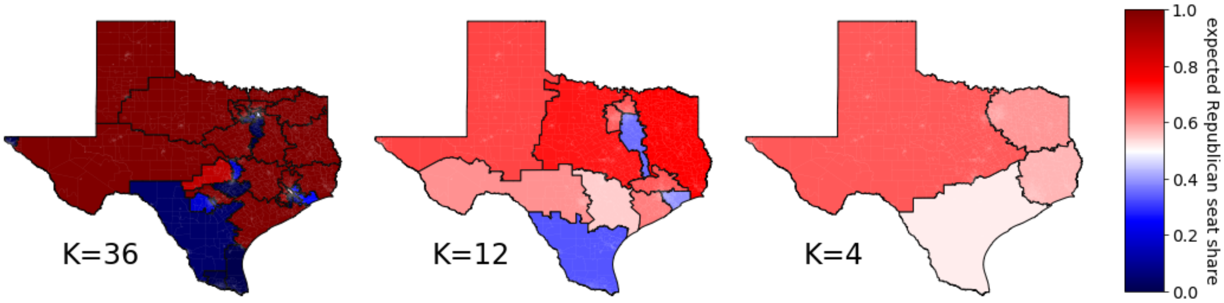

Figure 11 Example fair multi-member plans for Texas. Color denotes expected Republican seat share.

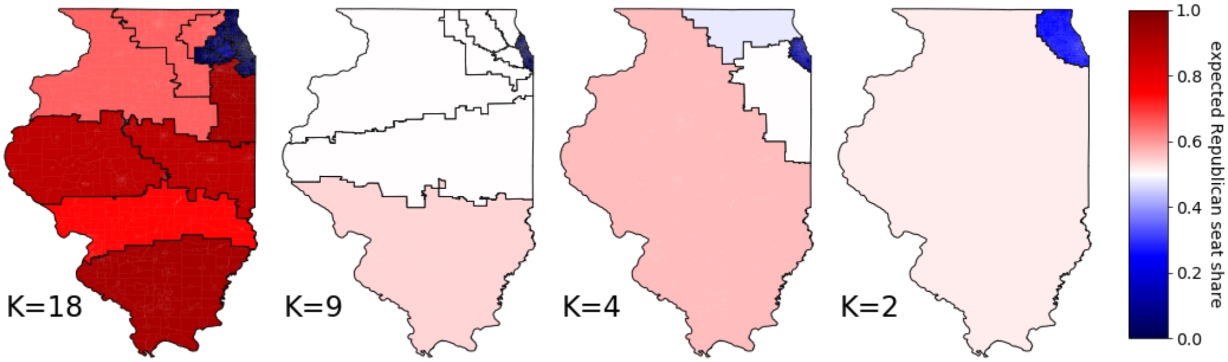

Figure 12 Example fair multi-member plans for Illinois. Color denotes expected Republican seat share.

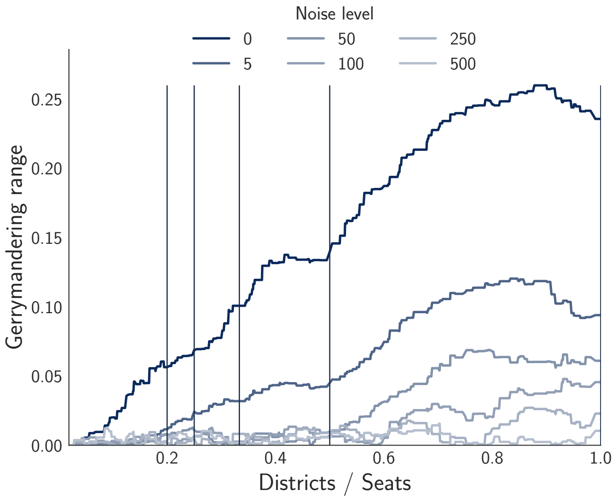

Figure 13 Replication of Figure 5 but with idiosyncratic noise drawn from a Normal distribution. Here, the noise level corresponds to the standard deviation σ. Results are qualitatively identical.

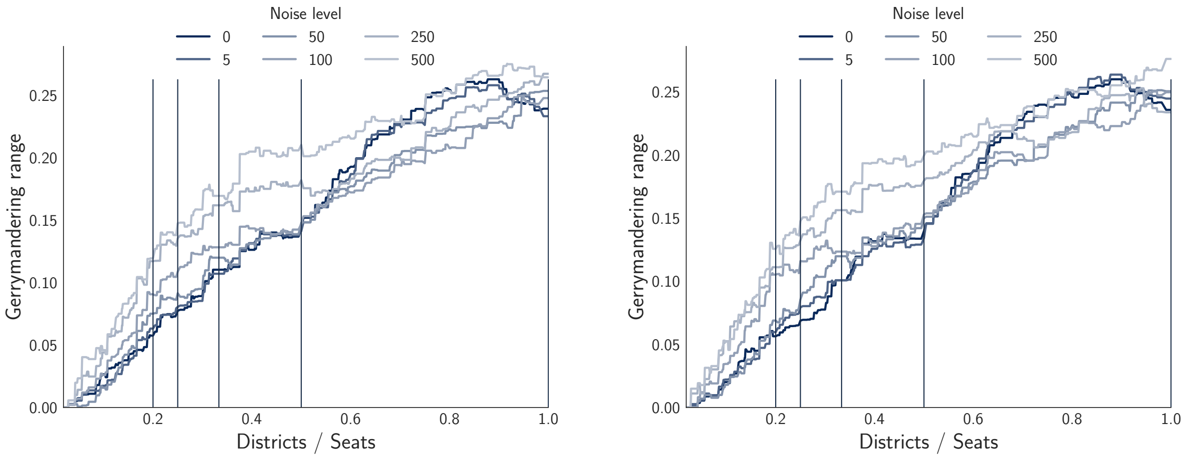

(a) Gumbel noise

(b) Normal noise

Figure 14 Replication of Figure 5 but with the optimizer having perfect knowledge of voter-noise realizations (which may also differ by map). Now, gerrymandering ability increases with knowable noise (since different maps may have different voter noise distributions, allowing for more extreme outcomes with noise). However, as before, gerrymandering range decreases with district size, with STV – even with an extreme case of optimizer knowledge, MMDs decrease their ability to construct disproportionate outcomes.

Appendix B: Results: Intra-party diversity

So far, we have studied how multi-member districts affect the balance of power between parties. However, one of the primary justifications of Ranked Choice Voting (either for SMDs or MMDs) is that it enables minor parties or ideologies to gain seats – it blunts the game-theoretic logic that tends to push winner-take-all democracies toward two parties. In this section, we analyze such claims by showing the effects of MMDs on intra-party diversity.

As in Section 4, the *solid coalitions* assumption of Proposition 1 does not hold. There, we analyzed crossover votes between parties; now, we analyze crossover votes *within each party* but between coalitions within that party; we must now generate intra-party rankings and simulate STV.[22] There is a further conceptual challenge: *coalitions* within parties are not well-defined, and identifying them with data is challenging; with many candidates, no two voters may share the same exact ballot preference order. In theory, results are often proven with regard to arbitrary-sized coalitions that approve the same candidates (cf. Skowron (2021)), but relating the guarantees back to coalitions in practice is difficult; it is unclear which coalitions proportionality should be defined with respect to, especially with limited ranked choice data, and how exactly to define proportionality for these groups.

We leverage additional political structure to tackle these challenges. (1) First, instead of considering arbitrary coalitions, we examine voter rankings emerging from differing assumptions on how voters value two dimensions: (approximate) political preference, and geographic preference. We choose these dimensions since single-member districts automatically impose one dimension (geographic) as more important, while in theory MMDs allow voters to prefer political connection over geographic proximity. (2) Second, we measure intra-party effects through two measures: the *diversity of the winning candidate set*, and the *diversity of the coalitions supporting each winner*. If the former increases, then that means more distinct within-party preferences are represented by the winning set. If the latter increases, that means each particular winner is accountable to a more diverse intra-party voter base. However, we do not claim that these dimensions are the most important – our simulation provides a method to understand how various dimensions trade off; insight for any given political setting requires careful consideration of the most important dimensions and their relationships. Rollout of MMDs might further affect partisan behavior and relationships, making it challenging to study such effects precisely before an implementation.

---

[22] In this section, we exclusively consider STV, as it allows voters within a party to prioritize different candidates, without risking their party overall representation by not approving a same-party candidate (as could happen with approval voting based methods like Thiele rules). Studying approval methods would require analyzing primaries, to select exactly N_k candidates from each party.

B.1. Methods and assumptions

To study the case when the solid coalitions assumption is broken, we first need to construct plausible intra-party voter rankings and candidate distributions; second, we need to simulate STV elections given a map. Methodological details are deferred to Appendix B.3. These methods are largely similar to those of Section 4.

Generating voters and candidates. For voters, we simply use the voters in each state in the voter file – subsampling 5,000 voters per district in each simulation. For replication purposes while preserving data privacy, in our code repository we provide a subsample of 50,000 voters, with noise added to the scores. We further generate one candidate per (party, census tract) combination for each simulation, with a maximum of 1,000 candidates per election.

Constructing intra-party rankings. We require further assumptions on how voters differentially rank candidates *within* their preferred party, to study how MMDs may change the characteristics of winners within a party. The key challenge is to develop a model for how a voter – given their characteristics – will vote given a menu of (hypothetical) candidates. Up to now, we have only assumed that voters approve all candidates of their party or rank them all above candidates of the other party.

We do so as follows, using the voter file described in Section 2.3. Recall that we have individual-level voter data in each tract, along with demographic information and ideological scores; in particular, we use the voter's geographic location (census tract of home address) and a univariate *partisan score* indicating the strength of their party affiliation (most Republican to most Democratic). We generate many synthetic candidates for each party in each district, with varying partisan scores; the distribution of candidates reflects that of the voters.

Finally, we assume that voters rank all candidates. They rank all same-party candidates over other-party candidates. Then, within each party, they rank candidates in order of the distance between their characteristics (either partisan scores or geographic location, in different simulations).[23] As a result, the *solid coalitions* assumption no longer holds within parties, where voter preferences may not be expressed neatly in terms of subparties.

[23] This approach is conceptually related to that of Becker et al. (2021), who use Ecological Inference (EI) to study racial groups' vote choices in primaries, to study how to draw SMDs such that a minority group's preferred candidate wins both a primary and the general election – with the insight that racial composition does not solely determine whether a district is effective for minorities, as within-party vote choice may not only depend on race. Our approach replaces EI with a calibrated, individual-level voter file; i.e., we assume that voters order candidates using their individual-level partisan scores and allow for the possibility that voters with different demographic characteristics may nevertheless vote similarly. For MMD analysis, both methods require extrapolating the scores to rankings, and thus require assumptions on how voters will vote for hypothetical candidates. Benade et al. (2021a) use a combination of historical voting behavior and ranked-choice model assumptions to generate such rankings, while we sample candidates and impose a spatial model between candidates and voters.

We note that while there has been much work on spatial voter models aiming to characterize such behavior based on the "distance" between the voter and each candidate (Adams et al. 2017, Tausanovitch and Warshaw 2018, Shor and Rogowski 2018, Jessee 2009), it is a hard challenge – it is not clear that voters behave according to such ideological spatial positioning, beyond the relative consistency of which party one votes for. Thus, our results should not be interpreted as what *would* necessarily happen with multi-member districts, but what intra-party coalitions *could or could not* be formed with MMDs given voter interest, and how these coalitions differ from those possible under SMDs.

Running STV elections. Given the candidates, sampled voters, and the voters' rankings over the candidates, we run fractional STV for each district in each given map. Fractional STV is STV as defined above. In each round, either a winner is selected if they have at least the Droop quota number of first-place votes, or a candidate with the least votes is eliminated. This candidate's votes are transferred, by eliminating this candidate from each voter's list. In fractional STV, votes are transferred as follows. Formally, suppose the winning candidate receives $v > Q$ first-place votes. Then, a fraction $\frac{v-Q}{v}$ of this candidate's votes is in excess of what is needed to win. Thus, each voter for this candidate has their weight multiplied by $\frac{v-Q}{v}$. For example, suppose the Droop quota was 5, and a candidate received 6.3 first-place votes (fractional votes are possible due to earlier round transfers). Then, each of their voters has their weight multiplied by $\frac{1.3}{6.3}$.

We simulate STV for the random, neutral maps calculated above, as we wish to study intra-party effects as distinct from partisan ones. Running these STV elections utilized over 60 CPU-weeks of compute, on top of the map generation discussed in Section 2.3. This runtime is for *given* maps, underscoring the necessity of Proposition 1 to study partisan seat shares without needing to simulate STV during the redistricting optimization process.

B.2. Results

Figure 15 contains our intra-party results, when we assume that within a party voters rank according to partisan score. First, Figure 15a illustrates how MMDs affect the diversity of the winning set within each party. For each map, we determine the set of winners for each party and then calculate the standard deviation of their partisan scores. For geographic diversity, we calculate the average Euclidean distance of each winner from the centroid of the winners' locations. We find that with STV and large districts, a more diverse set of winners emerges, within each party. The intuition for partisan score is simple and is similar to that regarding the diffusion of minority party voters across a state (the Massachusetts problem). When voters rank according to partisan score, then

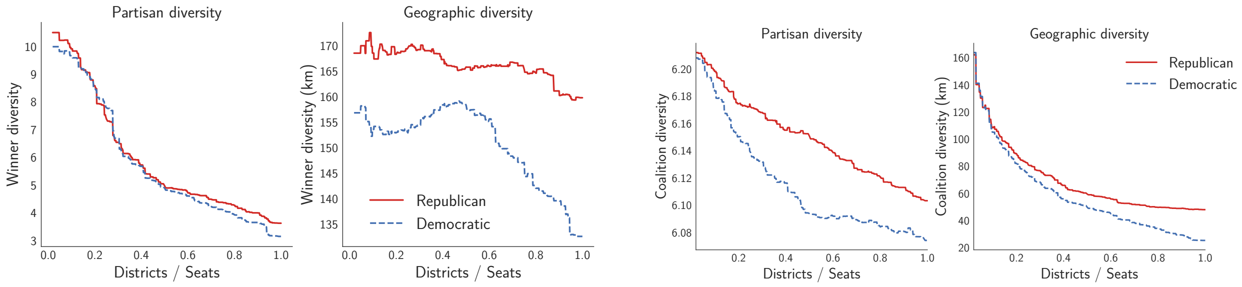

(a) Diversity of winning candidate set

(b) Diversity of coalitions supporting winners

Figure 15 How intra-party diversity measures vary with the number of districts. (a) The diversity of the winning candidate set. (b) The diversity of the *voters* who supported a given winning candidate, averaged by party. These results establish that, with STV, a more diverse set of winners can be elected, i.e., minority viewpoints *within* a party are supported. Simultaneously, each winning candidate draws support from a more diverse coalition of voters. Republicans (red line) tend to be more geographically spread out than Democrats (blue line), and so have a higher level of geographic diversity as measured at every district size. Republicans having higher partisan diversity coalitions may be related to the same phenomenon, since coalitions to gather winning votes may also be more spread out geographically and for partisan scores.

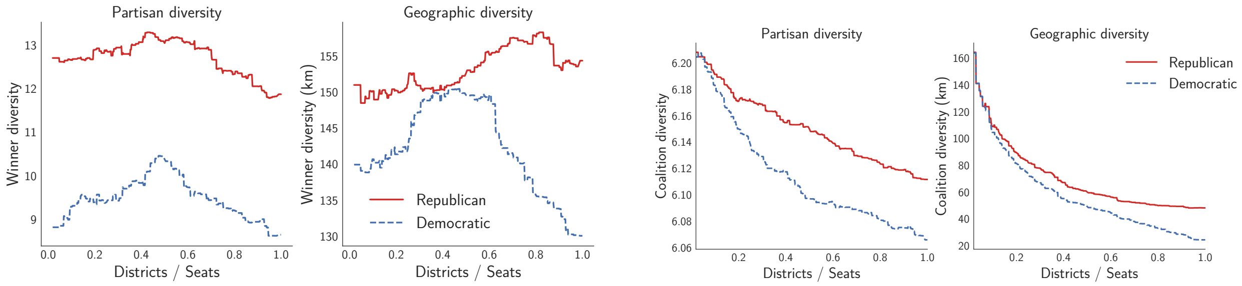

(a) Diversity of winning candidate set

(b) Diversity of coalitions supporting winners

Figure 16 Same as Figure 15, except voters now rank within party based on geographic distance.

similar voters across the state can pool their votes to ensure that a favored candidate wins. Using SMDs, these voters may be split across different districts, such that in each a different intra-party coalition chooses the winning candidate. Surprisingly, larger districts simultaneously increase the geographic diversity of winners, even when voters rank according to partisan score. These distance functions are defined in more detail in Appendix B.3.

Second, Figure 15b shows that with MMDs each winning candidate draws support from a more diverse coalition of voters. For each winner, we determine the voters who contributed votes in the STV round in which the candidate was elected, weighted by how many votes they provided (since we do fractional STV). For partisan score, we calculate the (weighted) standard deviation of the partisan scores of the voter coalition. For geographic diversity, we calculate the (weighted) average voter Euclidean distance from the district centroid. Finally, we average across winners within each party.

This result establishes that MMDs may come at a cost, in terms of the *geographic representation* aspect of our representative democracy: insofar as it is valuable for a representative to be accountable to a geographically cohesive set of voters (such as for providing constituent services or acquiring funding for projects), large MMDs weaken such ties; voter coalitions move from about 25 kilometers from the coalition center on average to almost 160 kilometers. Two- or three-member districts, however, come at a far smaller cost.

Appendix Figure 16 contains the same plots when voters' intra-party rank is according to geography. Perhaps as expected, under this assumption MMDs do not increase the diversity of the winners according to partisan scores – if voters do not band together based on partisan scores, then their mutually preferred candidates on this dimension do not win. However, the findings regarding *coalition* diversity remain virtually identical: even when voters rank within a party based on geographical distance to a candidate, each winner represents a far more geographically dispersed coalition than with SMDs.

While this finding may seem paradoxical, it can occur when voters for one party are not evenly distributed throughout the state – whereas voters in a (smaller) city may be large enough under SMDs to elect a preferred candidate, in MMDs they may be overruled by a larger group of same-party voters in another city. This finding adds a warning to the notion of proportionality—small, potentially intersectionally defined groups may be better served by small districts, even as theoretical proportionality guarantees a monotonic increase in district size. Most practically for the United States, more research is needed on the effect of MMDs on minority voters (cf. Benade et al. (2021a)).

Overall, our results caution in choosing too large a district size for MMDs, especially as our results in Section 3 suggest that most of the proportionality benefits can be achieved with two- or three-member districts—and results in this section suggest that such districts come at a smaller cost in terms of geographic cohesiveness. More broadly, this section gives an approach to evaluating the effects of voting rules and redistricting on intra-party outcomes, by generating various voter rankings from individual-level data and then studying the effects with respect to interpretable dimensions.

B.3. Intra-party method details

Here, we provide additional detail for Appendix B.

Voter multi-dimensional preferences. In the voter file provided to us, each voter is assigned scores on a variety of ideological and demographic dimensions, including party preference and preferences on specific issues. See the following blog post for additional details on these scores, under "Survey-based machine learning support scores": `https://predictwise.medium.com/what-all-we-did-for-the-2020-elections-78b66a44f651`. Here we use the *partisan score* value, which is an overall score from 0 to 100 designed to represent how likely a voter is to lean *Democrat* in the two-party vote; this score is derived from party registration, past primary voting data, demographic information, etc. As the post shows, this score correlates well with issue-preference scores. We threshold this value; above the threshold, a voter is declared a *Democrat* and below the threshold, a *Republican*.

As a second dimension, we use *geography* – based on the voter's home location; for computational simplicity and privacy reasons, we use the centroid of the census tract in which the voter's home location resides, as opposed to the latitude/longitude information of the voter.

Candidate generation. Next, we generate one candidate for each (party, census tract) combination as follows. (Results do not change with multiple candidates for each (party, census tract) combination). The home address of the candidate is taken to be the centroid of the census tract. The *partisan score* of the candidate is taken to be the median partisan score of voters in the corresponding party in that tract.

Voters and candidates in an election. Recall – from the map generation stage – that a given *map* is a partition of the census tracts into districts. Note that each voter and candidate belongs to a census tract, and thus a unique district in each map.

In a given election (as characterized by the map, social choice function, and number of winners in each district), each voter has available to them a subset of the candidates (those belonging to a census tract in the same district as the voter). The method to generate candidates might produce many candidates or voters in a given district (e.g., if the district makes up the entire state); for computational reasons, we sample at most 5,000 voters and 1,000 candidates from each district.

Turning voter preferences into rankings. In each simulated election, we need a ranking over candidates from each voter. These are generated as follows. First, as in the assumption of Proposition 1, we assume that each voter ranks their own party candidates over all other-party candidates. Within each party, voters rank candidates according to a *distance function*. As introduced in Appendix B, for some settings we assume that the distance function is first determined by the *partisan score* – voters rank candidates according to the absolute value of the difference of their partisan scores, with ties if any broken by geography. In other settings, we assume that the distance function is determined first by *geography*, with voters preferring closer candidates according to the Euclidean distance between the latitude/longitude associated with each voter and candidate (and ties broken by partisan score). Unlike reality, we assume that there is no *ballot exhaustion* – these simulated voters submit complete rankings.

Simulating STV elections. From the above steps, for each election in each district, we have a set of up to 5,000 voters and 1,000 candidates who have submitted rankings over the candidates. We run fractional STV for each such district.

This entire process took over 60 CPU-weeks to compute for all settings, with the largest computational load being generating the individual voter rankings and then running fractional STV.

Appendix C: Proofs

LEMMA 1. *Let* $y_R \in (0,1]$. *Consider*

$$n_R(y_R, \lambda_{PAV}) = \min_n \left[\arg\max_n \left[y_R \sum_{i=1}^{n} \lambda(i) + (1-y_R) \sum_{i=1}^{M-n} \lambda(i) \right] \right].$$

and let ℓ *be the unique integer such that*

$$y_R(M+1) - 1 \leq \ell < y_R(M+1)$$

Then, $n_R(y_R, \lambda_{PAV}) = \ell$*, for all* y_R, M.

We have that $n_R(y_R, \lambda_{PAV})$ is n such that:

$$\begin{aligned} n \geq 1 &\implies y_R\lambda(n) > (1-y_R)\lambda(M-n+1) \\ n < M &\implies y_R\lambda(n+1) \leq (1-y_R)\lambda(M-n) \end{aligned}$$

The first condition requires that choosing the nth candidate from party R is strictly (due to tie-breaking against R) more valuable than choosing the $(M-n+1)$th candidate from party D.

The second condition requires that choosing the $(n+1)$th candidate from party R is no more valuable than choosing the $(M-n)$th candidate from party D.

Rewriting the first condition, we have

$$\begin{aligned} n \geq 1 \implies \frac{y_R}{n} &> \frac{1-y_R}{M-n+1} \\ \iff \frac{y_R}{n} + \frac{y_R}{M-n+1} &> \frac{1}{M-n+1} \\ \iff \frac{y_R(M+1)}{n(M-n+1)} &> \frac{1}{M-n+1} \\ \iff \frac{y_R}{n} &> \frac{1}{(M+1)} \end{aligned}$$

Similarly, rewriting the second condition yields

$$n < M \implies \frac{y_R}{n+1} \leq \frac{1}{(M+1)}$$

Proposition 1 (Seat shares under STV) *Suppose – for a given district with* M *seats and* $V \geq M(M+1)$ *voters – that a fraction* y_p *of voters belong to each party* $p \in \{R, D\}$*, and that there are at least* M *candidates per party. Assume that each party's voters rank all same-party candidates above all other-party candidates, and ties are broken in party* D*'s favor.*

Let $n_R(y_R, STV)$ *be the number of winners belonging to party* R *using STV. Then, both* $n_R(y_R, STV)$ *and* $n_R(y_R, \lambda_{PAV})$ *are in* $\{\lfloor y_R M \rfloor, \lceil y_R M \rceil\}$*.*

The first part of the result is known from Tideman (1995) and Dummett (1984). For completeness, we provide a proof using our notation and in our exact setting.

For expositional simplicity, we assume that $V \bmod (M+1) = 0$, i.e., the number of votes is evenly divisible by one plus the number of winners, though the same arguments extend. Then, let $Q = \lfloor \frac{V}{M+1} \rfloor + 1 = \frac{V}{M+1} + 1$ be the Droop quota.

To understand the intuition, note that the given seat shares would immediately hold if each party had a coordinator who could optimally decide how voters of that party rank candidates within the party. In that case, the coordinator would ensure that as many candidates as possible have first-place votes equal to the Droop quota, with no first-place votes going to candidates who will be eliminated. The resulting seat shares follow, given tie-breaking against party R (and, if needed, the surplus vote transfer procedure for the additional single vote needed in the Droop quota).

However, this argument is not sufficient because, in principle, candidates could receive meaningful votes from voters of the other party. Further, without such a coordinator, a sub-optimal arrangement of votes could potentially lead to the elimination of candidates who would be elected with such a coordinator. (The former reason is fundamental, and is why the exact formula does not hold in general without parties or even for more than two parties without further assumptions; the latter possibility is exactly the issue that STV surplus vote transfers are designed to avoid, and eliminating it is bookkeeping). The proof centers around eliminating these two possibilities.

The key step in the proof is noting that, under the assumptions, votes can only be transferred from a candidate of one party to a candidate of the other party (either after a candidate is eliminated or selected as a winner) if candidates from the sending party have been exhausted, as the sending party voters rank all other party members after all their own party members.

Thus, for any set of voter rankings under the assumptions, the per-party seat share remains the same under re-arrangements of how each voter ranks members of the *other* party (at the point that such rearrangements matter for who is elected, only one party's candidates are left, and so the partisan seat share does not change). Thus, without loss of generality, for the rest of the proof we assume that all voters within a party share the same ordering for candidates of the *other* party.

Recall that the Droop quota is designed such that exactly M total candidates meet the Droop quota across rounds, if all voters submit full rankings and $V \geq M(M+1)$. (For simplicity and without loss of generality, we assume that the election continues even when the number of candidates remaining equals the number of seats to be filled, so that they each meet the quota after transfers).

More than M candidates reaching the quota would require at least $(M+1)Q = V + M + 1 > V$ total votes (as votes necessary to reach the quota are never transferred).

Similarly, we have that at least M candidates reach the Droop quota after transfers. There are enough votes, optimally spread, for enough candidates to reach the quota:

$$\begin{aligned} MQ \leq V &\iff M\left[\frac{V}{M+1}\right] + M \leq V \\ &\iff \left[\frac{M}{M+1} - \frac{M+1}{M+1}\right] V \leq -M \\ &\iff -\left[\frac{1}{M+1}\right] V \leq -M \\ &\iff V \geq M(M+1) \end{aligned}$$

Note that the above argument holds at every round. Once we've elected W candidates, we need to elect $M - W$ more, and there are $V - WQ$ first-place votes left:

$$(M - W)Q = MQ - WQ \leq V - WQ$$

and so by the pigeonhole principle, when there are $M - W$ candidates left (possibly after eliminating some), there is at least one candidate with at least as many votes as the Droop quota.

The above facts establish that we can carry out the initial rounds of STV separately for each party, until in each party there are either no candidates remaining or there is one remaining candidate with less than Q votes: no winners have been selected with votes transferred across parties up to this stage; as long as we only elect candidates with votes at or above the Droop quota, we do not mistakenly elect any candidate separately that we would not have together; and if this stage is ever reached, the *identities* of the elected or eliminated candidates do not matter, since up to now there is a conservation of votes by party and so the per-party counts remain the same.

Now, consider each party p, and suppose in the current round that W_p candidates have been elected for the party. Recall that y_pV is the number of voters for party p.

Further suppose that there are $m_p \geq 1$ candidates left for the party, and they collectively have $q_p = y_pV - W_pQ$ first-place votes (either original first-place votes, or votes after transfers) among them. We repeat the above argument. If $q_p \geq Qm_p$, then by the pigeonhole principle, there exists at least one candidate of the party such that the votes for that candidate meet the Droop quota Q. Then, that candidate is declared a winner, and its surplus votes are transferred to other candidates of the same party. Then, iterate with $W_p + 1$, $m_p - 1$, and $q_p - Q$. Otherwise, eliminate the candidate in the party with the least number of votes, and iterate with $m_p - 1$ and q_p until $m_p = 0$ or $m_p = 1$ with the remaining candidate having less than Q votes.

Suppose at the end of these separate processes, W_p candidates from each party p have been elected. We know that there is at most one candidate from each party p remaining, with $y_pV - W_pQ$ votes.

Then

$$\begin{aligned}
n_p(y_p, STV) \geq W_p &\geq \lfloor \frac{y_pV}{Q} \rfloor \\
&= \lfloor \frac{y_pV}{\frac{V}{M+1} + 1} \rfloor \\
&= \lfloor \frac{y_p(M+1)}{1 + \frac{M+1}{V}} \rfloor \\
&\geq \lfloor \frac{y_p(M+1)}{1 + \frac{M+1}{M(M+1)}} \rfloor & V \geq M(M+1) \\
&= \lfloor y_pM \rfloor
\end{aligned}$$

Since this holds for both parties simultaneously, we have

$$\begin{aligned}\lfloor y_R M \rfloor &\leq n_R(y_R, STV) \\ &\leq M - \lfloor M - M y_R \rfloor = \lceil y_R M \rceil.\end{aligned}$$

Whether $n_R(y_R, STV)$ is at the ceiling or floor depends on whether it has the majority of the remaining votes when each party has at most one candidate remaining.

The proof is finished by applying Lemma 1.

Appendix D: Empirical methodological detail

D.1. Map Generation

All details of the original Stochastic Hierarchical Partitioning (SHP) algorithm can be found in the paper by Gurnee and Shmoys (2021) and associated appendices; we further use their same geographic and electoral data (see Appendix Table 2 of Gurnee and Shmoys (2021)). Here, we discuss the relevant algorithmic parameters and differences from that algorithm.

Multi-member. The main algorithmic difference from the original SHP algorithm is that we adapted ours to generate multi-member districts. To do this, instead of parameterizing a sample tree node by a region R and total number of seats s, we needed to also specify the number of districts that node contains. This is required because at an intermediate node, the number of districts is not immediately derivable from the total number of seats in that node because of ambiguity in the number of N_k versus $N_k + 1$ sized districts. Therefore, the number of seats is used just to balance population, and the number of districts is used for all other tree operations (sampling valid splits, maintaining balance, etc.).

Ensembles. For each pair (state, K), with $K \in \{2, 3, \ldots, 10, 12, \ldots, 20, 23, \ldots 53\} \cup \{N\}$ and $K \leq N$, we sampled the root node $\left(\frac{1000}{K}\right)^{1.2}$ times and each internal node $\left(\frac{300}{K}\right)^{0.5}$ times. These constants were chosen to balance computational cost and optimization quality. We used random-iterative center selection with Voronoi-weighted capacity matching to sample region centers and sizes. All districts are population balanced within a 1% tolerance. Each of these ensembles was then scored, optimized, and subsampled to create a final distribution over partisan outcomes.